

NERSC Metrics



Discussion with Office of Science NERSC
Allocation Managers
October 2022

Operational Metrics

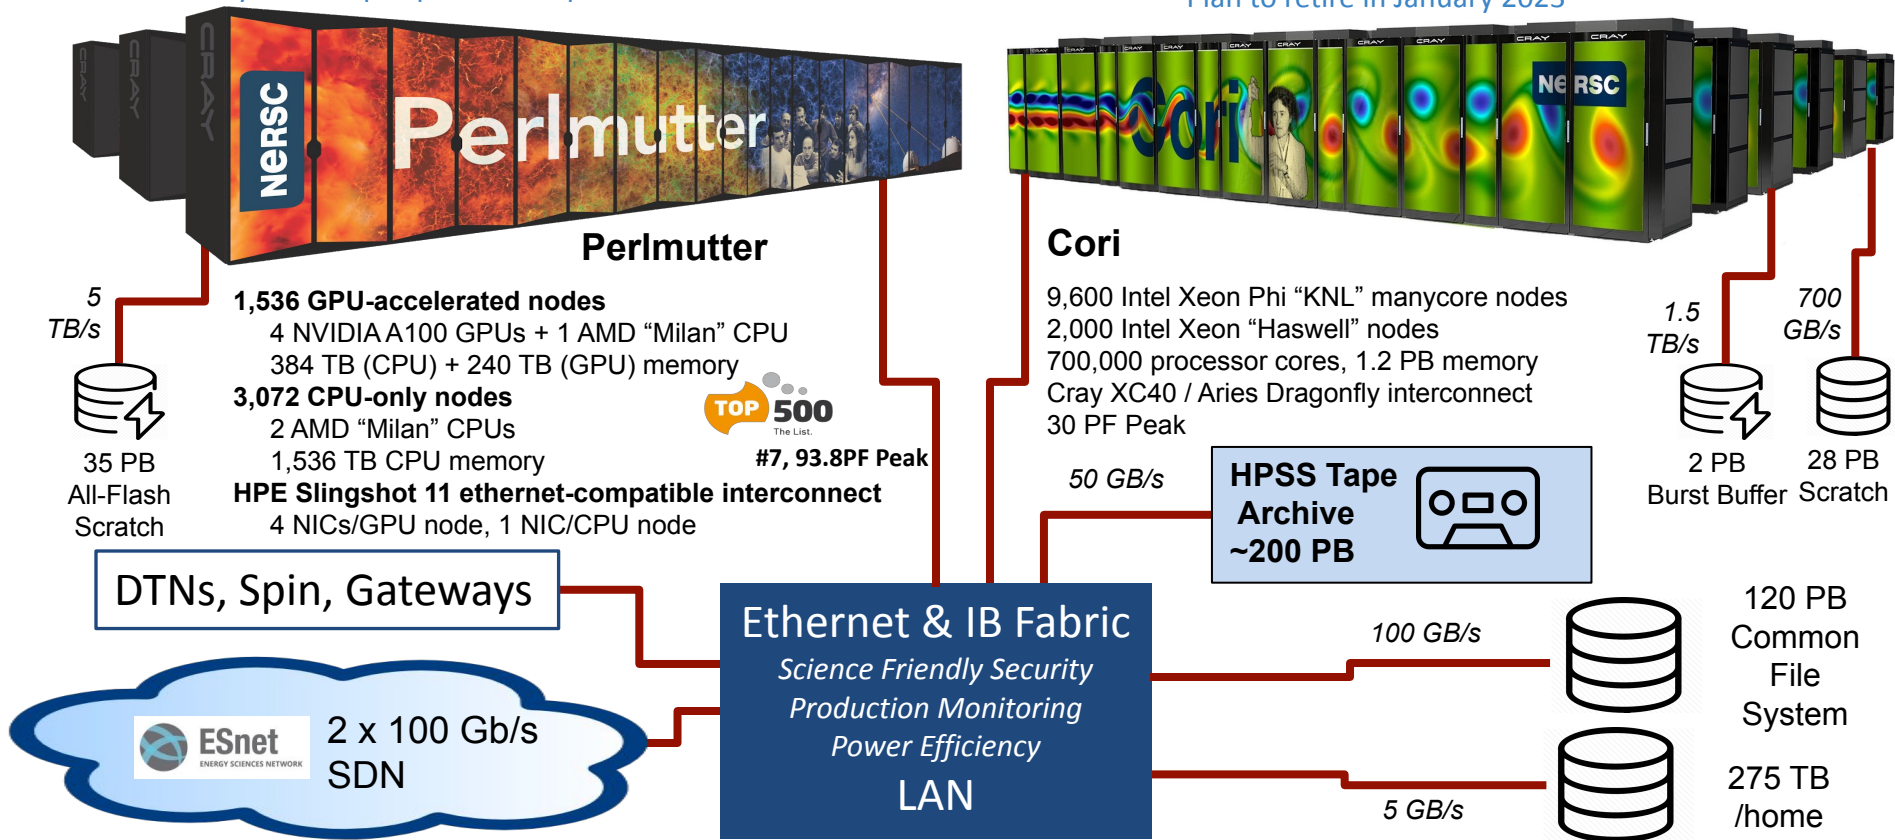
Operational metrics are indicators that measure an organization's performance. These numbers provide a snapshot of key processes like services delivered. Operational metrics tracking shows how well the organization performs these processes. The data may measure efficiency, productivity or quality.




NERSC Systems Fall 2022

Still in Early Science pre-production phase

Plan to retire in January 2023




NERSC by the Numbers




NERSC by the Numbers

Top Science Disciplines

(By computational hours used)



Breakdown of Compute Used by DOE Program



>2,000
Scientific Journal Articles
per Year

NERSC Status


- Cori is running productively as usual

- Scheduled availability > 99%
- Overall availability > 98%
- Utilization = 93.9%


- Perlmutter

- GPU-accelerated and CPU-only nodes are being integrated into single system
- Network upgrade to “Slingshot 11” has been completed
- Scratch file system “Lustre” upgrade is nearing completion
- Heavy usage continues during “free early science” period
- Charging expected to start soon (assuming scratch file system remains stable)


- System has been available to users during upgrades and integration**



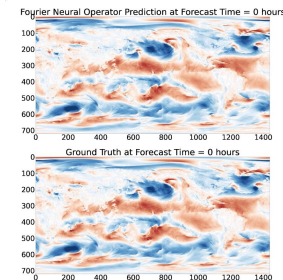
3D computed structural modeling of protein interactions is made possible by deep learning and evolutionary analysis. Credit: University of Washington (Ian Haydon)




World's first exa-op electronic structure calculation using mixed precision arithmetic and Perlmutter CPU tensor cores. SARS-CoV-19 spike protein shown.




DESI spectral extraction




Data-Driven Forecasting



Advances in understanding heat transport via turbulence in NSTX-U




Open Catalyst Project



Calculations support interpretation of Fermilab Muon g-2 CP violation experiment




System availability excluding scheduled outages.



Cori: 98.6%
HPSS: 100%
CFS: 100%

Month	Cori (%)	HPSS (%)	CFS (%)
Oct 21	99.4	100	100
Nov 21	99.8	100	100
Dec 21	99.4	100	100
Jan 22	99.6	100	100
Feb 22	100	100	100
Mar 22	98.1	100	100
Apr 22	97.6	100	100
May 22	94.2	100	100
Jun 22	98.7	100	100
Jul 22	99.8	100	100
Aug 22	97.7	100	100
Sep 22	99.3	100	100


System overall availability



Cori: 97.8%
 HPSS: 99.8%
 CFS: 100%

Month	Cori (%)	HPSS (%)	CFS (%)
Oct 21	97.5	100	100
Nov 21	99.8	100	100
Dec 21	97.7	99.8	100
Jan 22	98.1	99.7	100
Feb 22	100	100	100
Mar 22	98.1	100	100
Apr 22	97.6	99.5	100
May 22	94.6	99.4	100
Jun 22	98.7	99.7	100
Jul 22	98.0	99.7	100
Aug 22	95.7	100	100
Sep 22	97.6	100	100


Percent of available time used by jobs on Cori



Cori: 93.8%

Month	Cori (%)
Oct 21	93.8
Nov 21	93.1
Dec 21	94.0
Jan 22	93.9
Feb 22	93.6
Mar 22	93.1
Apr 22	93.9
May 22	94.8
Jun 22	90.5
Jul 22	95.6
Aug 22	94.0
Sep 22	95.1

Percent of Time Used by Jobs >1,024 Nodes on KNL




Cori: 35.5%

Month	Capability Metric (%)	Target (%)
Oct 21	20.9	25
Nov 21	25.7	25
Dec 21	38.7	25
Jan 22	32.7	25
Feb 22	37.5	25
Mar 22	32.7	25
Apr 22	31.2	25
May 22	41.2	25
Jun 22	41.7	25
Jul 22	50.3	25
Aug 22	39.0	25
Sep 22	34.7	25



Queue Wait Times


Average Regular Queue Wait Times




Month	Cori KNL (hours)	Cori Haswell (hours)
Oct 21	4.7	14.7
Nov 21	9.4	11.7
Dec 21	22.2	43.5
Jan 22	55.3	54.9
Feb 22	24.1	53.1
Mar 22	43.1	48.9
Apr 22	29.0	66.2
May 22	32.49	293.76
Jun 22	53.03	59.66
Jul 22	38.12	177.80
Aug 22	79.44	105.34
Sep 22	47.37	101.84

Consulting and Support

Number of New Tickets and Backlog of Open Tickets



3-Day Expectations Set Metric




8,543 User Tickets in last 12 months

Only 36 (0.4%) “not satisfied” on follow up surveys



Perlmutter GPU Nodes (Phase 1)

- Early users on the system since July 2021 & all user accounts created in May 2022
- 494 projects and 1,672 users have run batch jobs on the system in 2022
- System has been highly productive across all DOE Office of Science programs
- 5 Million GPU Node Hours delivered to science in 2022




Perlmutter GPU Node Usage 2022

GPU Job Backlog: 81 days


Jobs Have Run at All Scales on GPU Nodes

- Full system Gordon Bell runs
 - 928 node WarpX
 - 952 node FourCastNet
 - 1,024 node PASTIS (ExaBiome)
- ECP Application Development & Testing
 - 394 jobs used more than 1,000 nodes
 - 43 teams*



CPU-Only Nodes (Phase 2)


- CPU-only nodes started being integrated into the system in Spring and were opened to users in May
 - 436 projects and 1,059 users have run batch jobs on the CPU nodes
 - Half the nodes are currently in the system; the remaining half will be added this year
- Easy transition for Cori workloads
- 4.1 Million CPU Node Hours delivered (free) to science




Perlmutter CPU Node Usage 2022

CPU Job Backlog: 21 days

GPU Node Utilization Is Extremely High





CPU Node Utilization Has Been Extremely High




User Rate NERSC Highly on Survey

Overall Satisfaction



Average scores for all questions



Comments

What Does NERSC Do Well?


- Consulting, account support, and training
- Providing high-quality, usable resources
- Documentation
- Software
- Communication with users.

What Can NERSC Do Better?

- Reduce NERSC system maintenance periods
- Reduce queue wait times
- Fewer storage system interruptions
- Improve documentation
- Provide more CPU resources.



Review of 2022 Allocation Year to Date



2022 Allocations

- NERSC provided both CPU and GPU time to DOE Mission Science for Allocation Year 2022 and NERSC is on pace to exceed both targets
 - CPU Node Hours: Available: 24,650,000, Distributed by Program Managers: 22,237,619
 - 17,500,000 (raw) CPU Node hours have been used to date
 - GPU Node Hours: Available: 3,400,000, Distributed by Program Managers: 3,044,028
 - 2,700,000 (raw) GPU Nodes hours have been used to date
- Six months of charging on Perlmutter was built into NERSC's model
 - Supply chain and other issues have delayed the start of charging
 - We believe the system will be ready to begin charging soon, which will give priority on Perlmutter to projects that got allocations from the Office of Science Program Managers
 - No plan to reduce GPU allocations or redistribute time




GPU Usage vs. Allocation

The correlation between allocation and usage during the free period is very loose.


There are various options on how to (re)allocate time for the remainder of the year.

The simplest leaves allocations as they are, which gives projects with larger allocations priority and accommodates other through backfill/overrun.


GPU Machine Hours vs. GPU Compute Allocation



Mission Science Hours - CPU



Mission Science Hours - GPU



2022-25 Allocations & Projections

DOE Mission Science Allocations			
Year	NERSC Hours (2021 Units)	CPU Node Hours (2022 Units)	GPU Node Hours
2021	6,834,000,000	16,000,000	-
2022	9,860,000,000	24,650,000	3,400,000
2023		15,400,000	6,850,000
2024		15,400,000	6,850,000
2025		15,400,000	6,850,000

We are assuming that the share of the total allocation available to each Program Office remains the same during this period.



Questions?

